arXiv:1605.07103v2 [math.SP] 3 Jun 2016

Decomposing Real Square Matrices via Unitary Diagonalization

Théo Trouillon*, Christopher R. Dance†, Éric Gaussier‡and

Guillaume Bouchard§

Abstract

Diagonalization, or eigenvalue decomposition, is very useful in many areas of applied mathematics, including signal processing and quantum physics. Matrix decomposition is also a useful tool for approximating matrices as the product of a matrix and its transpose, which relates to unitary diagonalization. As stated by the spectral theorem, *only* normal matrices are unitarily diagonalizable. However we show that *all* real square matrices are the real part of some unitarily diagonalizable matrix.

1 Introduction

Matrix decomposition is a very well studied field of linear algebra, in particular, eigenvalue decomposition has applications in many scientific areas. A specific case arises when the eigenvectors form an orthonormal basis. In other words, the matrix formed by the column eigenvectors is orthogonal. In this case we speak of orthogonal diagonalization.

Definition 1. *A real square matrix $A \in \mathbb{R}^{n\times n}$ is orthogonally diagonalizable if it can be written $A = S\Lambda S^\top$, with $S, \Lambda \in \mathbb{R}^{n\times n}$, Λ diagonal, and S orthogonal: $SS^\top = S^\top S = I$.*

Decompositions of n-by-n real matrices A that can be expressed as $A = S\Lambda S^\top$ are especially relevant when rows and columns of a square matrix represent the same objects of some underlying problem. Examples of motivating applications include spectral analysis of graphs [Cvetković et al., 1997], and decomposition of covariance matrices such as principal component analysis [Jolliffe, 1986].

*Xerox Research Centre Europe, Université Grenoble Alpes, theo.trouillon@xrce.xerox.com

†Xerox Research Centre Europe, chris.dance@xrce.xerox.com

‡Université Grenoble Alpes, eric.gaussier@imag.fr

§University College London, g.bouchard@cs.ucl.ac.uk

The spectral theorem tells us that a matrix is orthogonally diagonalizable if and only if it is symmetric. However in many cases asymmetric square matrices arise, and yet it would still be useful to have a symmetric decomposition of these matrices. This is frequently the case in machine learning problems, for example: learning vectorial representations of words from co-occurence matrices [Pennington et al., 2014], and of entities from knowledge graphs [Nickel et al., 2015]. In these applications, row i of the matrix represent the same object as column i. Therefore it would be natural to have unique representations, or embeddings, s_i for each object. But representing element A_{ij} of the matrix by $s_i^\top \Lambda s_j$ would yield a symmetric matrix A. To avoid this problem, practitioners either use sophisticated transformations of the original matrix to symmetrize it, or they consider more complicated decompositions, or they simply use non-symmetric decompositions like singular value decomposition.

We introduce a new decomposition using unitary diagonalization, the generalization of orthogonal diagonalization to the complex field. This allows decomposition of *arbitrary* real square matrices with unique representations of rows and columns. Let us first recall some notions of complex linear algebra.

Let $\overline{A} \in \mathbb{C}^{m \times n}$ denote the complex conjugate of the complex matrix A. We shall write $A^* \in \mathbb{C}^{n \times m}$ the conjugate-transpose $A^* = (\overline{A})^\top = \overline{A^\top}$. The conjugate transpose is also often written $A^\dagger$.

Definition 2. *A complex square matrix $A \in \mathbb{C}^{n \times n}$ is unitarily diagonalizable if it can be written as $A = S\Lambda S^*$, with $S, \Lambda \in \mathbb{C}^{n \times n}$, Λ diagonal, and S unitary: $SS^* = S^*S = I$.*

Definition 3. *A complex square matrix A is normal if it commutes with its conjugate-transpose $AA^* = A^*A$.*

We can now state the less well-known version of the spectral theorem that applies to the complex domain.

Theorem 1 (Spectral Theorem [Axler, 1997])**.** *Suppose A a complex square matrix. Then A is unitarily diagonalizable if and only if A is normal.*

Among normal matrices, there are some purely real matrices that are not symmetric, such as the skew-symmetric matrices, and thus can be diagonalized in the complex domain. As we only focus on *real* square matrices in this work, let us summarize all the cases where A is real square and $A = S\Lambda S^*$, with $S, \Lambda \in \mathbb{C}^{n \times n}$, Λ is diagonal and S is unitary:

- A is symmetric if and only if A is orthogonally diagonalizable, with S and Λ are purely real.
- A is normal and non-symmetric if and only if Λ is unitarily diagonalizable and S and Λ are *not* both purely real.
- A is not normal if and only if A is not unitarily diagonalizable.

In the following, we generalize all three cases by showing that, for any $A \in \mathbb{R}^{n\times n}$, there exists a unitary diagonalization in the complex domain, of which the real part equals A:

$$A = \mathrm{Re}(S\Lambda S^*).$$

In other words, the unitary diagonalization is projected on the real subspace.

2 Real Square Matrices

Theorem 2. *Suppose $A \in \mathbb{R}^{n\times n}$ is a real square matrix. Then there exists a normal matrix $X \in \mathbb{C}^{n\times n}$ such that* $\mathrm{Re}(X) = A$.

Proof. Let form the complex square matrix $X \triangleq A + iA^\top$, and derive X^*:

$$X^* = A^\top - iA = -i(iA^\top + A) = -iX.$$

Therefore X is normal:

$$XX^* = X(-iX) = (-iX)X = X^*X.$$

□

Remark that there also exists a normal matrix $X = A^\top + iA$ such that $\mathrm{Im}(X) = A$.

Following Theorem 1 and Theorem 2, any real square matrix can be written as the real part of a diagonal matrix through a unitary change of basis.

Corollary 1. *Suppose $A \in \mathbb{R}^{n\times n}$ is a real square matrix. Then there exist $S, \Lambda \in \mathbb{C}^{n\times n}$, where S is unitary, and Λ diagonal, such that $A = \mathrm{Re}(S\Lambda S^*)$.*

Proof. From Theorem 2, we can write $A = \mathrm{Re}(X)$, where X is a normal matrix, and from Theorem 1, X is unitarily diagonalizable. □

Let us discuss the rank of such diagonalizations. First, we recall the definition of the rank of a matrix.

Definition 4. $\mathrm{rank}(A) = k$ *with A an m-by-n matrix over an arbitrary field F, if A has exactly k linearly independent columns.*

Also note that if A is diagonalizable $A = S\Lambda S^{-1}$ with $\mathrm{rank}(A) = k$, Λ has k non-zero diagonal entries for some Λ diagonal and some matrix S invertible. From this it is easy to derive a known additive property of the rank: $\mathrm{rank}(B + C) \leq \mathrm{rank}(B) + \mathrm{rank}(C)$, with B, C m-by-n matrices over an arbitrary field F.

We now show that any rank-k real square matrix can be reconstructed from a $2k$-dimensional unitary diagonalization.

Corollary 2. *Suppose $A \in \mathbb{R}^{n\times n}$ arbitrary, and $rank(A) = k$. Then there exist $S \in \mathbb{C}^{n\times 2k}$, $\Lambda \in \mathbb{C}^{2k\times 2k}$, where the columns of S form an orthonormal basis of $\mathbb{C}^{2k}$, and Λ diagonal, such that $A = \mathrm{Re}(S\Lambda S^*)$.*

Proof. Consider the complex square matrix $X \triangleq A+iA^\top$. We have $\mathrm{rank}(iA^\top) = \mathrm{rank}(A^\top) = \mathrm{rank}(A) = k$.

Thus $\mathrm{rank}(X) \leq \mathrm{rank}(A) + \mathrm{rank}(iA^\top) = 2k$.

The proof of Theorem 2 shows that X is normal, thus: $X = S\Lambda S^*$ with $S \in \mathbb{C}^{n\times 2k}$, $\Lambda \in \mathbb{C}^{2k\times 2k}$ where the columns of S form an orthonormal basis of $\mathbb{C}^{2k}$, and Λ diagonal.

□

Given that such decomposition always exists in dimension n (Theorem 2), this upper bound is not relevant when $rank(A) \geq \frac{n}{2}$.

3 Low sign-rank Matrices

We interest ourselves here in square sign matrices, $Y \in \{-1, 1\}^{n\times n}$, and how they can be reconstructed using the sign function of a real matrix $X \in \mathbb{R}^{n\times n}$: $Y = \mathrm{sign}(X)$, where $\mathrm{sign}(X)_{ij} = \mathrm{sign}(x_{ij})$. There are many such matrices X. More precisely, all such matrices X constitute a different orthant of $\mathbb{R}^{n\times n}$ for each sign matrix Y.

The use of the sign function is here not arbitrary, as it maps to a known complexity measure for sign matrices: the sign-rank [Linial et al., 2007].

Definition 5. *The sign-rank $\mathrm{rank}_\pm(Y)$ of an m-by-n sign matrix Y, is the smallest rank among all the m-by-n real matrices that have the same sign-pattern as Y. Formally:*

$$\mathrm{rank}_\pm(Y) = \min_{A\in\mathbb{R}^{m\times n}} \{\mathrm{rank}(A) | \mathrm{sign}(A) = Y\}.$$

Using Corollary 2, we can now show that any square sign matrix of sign-rank k can be reconstructed from a rank-$2k$ unitary diagonalization.

Corollary 3. *Suppose $Y \in \{-1, 1\}^{n\times n}$, $\mathrm{rank}_\pm(Y) = k$. Then there exists $S \in \mathbb{C}^{n\times 2k}$, $\Lambda \in \mathbb{C}^{2k\times 2k}$ where the columns of S form an orthonormal basis of $\mathbb{C}^{2k}$, and Λ diagonal, such that $Y = \mathrm{sign}(\mathrm{Re}(S\Lambda S^*))$.*

Proof. By definition, if $\mathrm{rank}_\pm(Y) = k$, there exists a real square matrix A such that $\mathrm{rank}(A) = k$ and $\mathrm{sign}(A) = Y$. From Corollary 2, $A = \mathrm{Re}(S\Lambda S^*)$ where $S \in \mathbb{C}^{n\times 2k}$, $\Lambda \in \mathbb{C}^{2k\times 2k}$ where the columns of S form an orthonormal basis of $\mathbb{C}^{2k}$, and Λ diagonal. □

To the best of our knowledge, previous attempts to approximate the sign-rank did not use complex numbers. They showed the existence of compact factorizations under conditions on the sign matrix [Nickel et al., 2014], or only in specific cases [Bouchard et al., 2015]. This results shows that if a sign matrix has sign-rank k, then it can be reconstructed by a $2k$-dimensional unitary diagonalization in the complex space.

Example

Consider the following 2-by-2 sign matrix:

$$Y = \begin{bmatrix} -1 & -1 \\ 1 & 1 \end{bmatrix}$$

Not only is this matrix not normal, but one can also easily check that there is no *real* normal 2-by-2 matrix that has the same sign-pattern as Y. Clearly, Y is a rank-1 matrix since its columns are linearly dependent, hence its sign-rank is also 1. From Corollary 3, we know that there is a normal matrix of which the real part has the same sign-pattern as Y, and of which the rank is less than or equal to 2.

However, there is no such rank-1 unitary diagonalization of Y. Otherwise we could find a 2-by-2 complex matrix X such that $\mathrm{Re}(x_{11}) < 0$ and $\mathrm{Re}(x_{22}) > 0$, where $x_{11} = s_1\lambda\bar{s}_1 = \lambda|s_1|^2$, $x_{22} = s_2\lambda\bar{s}_2 = \lambda|s_2|^2$, $s \in \mathbb{C}^2, \lambda \in \mathbb{C}$. This is obviously unsatisfiable. This example generalizes to any n-by-n square sign matrix that only has -1 on its first row and is hence rank-1, the same argument holds considering $\mathrm{Re}(x_{11}) < 0$ and $\mathrm{Re}(x_{nn}) > 0$.

This shows that the upper bound on the rank of the unitary diagonalization showed in Corollaries 2 and 3 is strictly greater than k, the rank or sign-rank, of the decomposed matrix.

Though this decomposition is clearly not unique, it can capture meaningful representations of the input matrix in a low-rank decomposition setting. Trouillon et al. [2016] show state-of-the-art results by jointly decomposing a set of square sign matrices representing relational data.

References

Sheldon Jay Axler. *Linear algebra done right*, volume 2. Springer, 1997.

Guillaume Bouchard, Sameer Singh, and Théo Trouillon. On approximate reasoning capabilities of low-rank vector spaces. *AAAI Spring Syposium on Knowledge Representation and Reasoning (KRR): Integrating Symbolic and Neural Approaches*, 2015.

Dragoš M. Cvetković, Peter Rowlinson, and Slobodan Simic. *Eigenspaces of graphs*. Number 66. Cambridge University Press, 1997.

Ian T. Jolliffe. *Principal Component Analysis*. Springer Verlag, 1986.

Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.

Maximilian Nickel, Xueyan Jiang, and Volker Tresp. Reducing the rank in relational factorization models by including observable patterns. In *Advances in Neural Information Processing Systems*, pages 1179–1187, 2014.

Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs: From multi-relational link prediction to automated knowledge graph construction. *arXiv preprint arXiv:1503.00759*, 2015.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543, 2014.

Théo Trouillon, Johannes Welbl, Guillaume Bouchard, Sebastian Riedel, and Eric Gaussier. Complex embeddings for simple link prediction. In *Proceedings of the 33rd International Conference on Machine Learning (ICML)*, 2016.